A Real-Time Pedestrian Detector using Deep Learning for Human-Aware Navigation

David Ribeiro, André Mateus, Jacinto C. Nascimento, and Pedro Miraldo

arXiv:1607.04441v1 [cs.RO] 15 Jul 2016

***Abstract*—This paper addresses the problem of Human-Aware Navigation (HAN), using multi camera sensors to implement a vision-based person tracking system. The main contributions of this paper are a novel and real-time Deep Learning person detection and a standardization of personal space, that can be used with any path planer. In the first stage of the approach, we propose to cascade the Aggregate Channel Features (ACF) detector with a deep Convolutional Neural Network (CNN) to achieve fast and accurate Pedestrian Detection (PD). For the personal space definition (that can be defined as constraints associated with the robot's motion), we used a mixture of asymmetric Gaussian functions, to define the cost functions associated to each constraint. Both methods were evaluated individually. The final solution (including both the proposed pedestrian detection and the personal space constraints) was tested in a typical domestic indoor scenario, in four distinct experiments. The results show that the robot is able to cope with human-aware constraints, defined after common proxemics rules.**

I. Introduction

For robots to interact naturally with humans in their social environments, they must have the ability to plan their motion accounting for the typical social norms. In this paper, we address the robot navigation in the presence of humans, resorting on multi cameras (static outside and/or on-board cameras) for the vision-based person tracking system.

In the last few years, robotics has become focused on Human-Robot Interaction and on its role in social environments. When people think of a robot interacting with a person, what comes to mind is a robot that can speak with her or hand over some object. However, the motion itself is of great importance in a social context (e.g., when a robot is requested to fetch an item), or simply when a normal navigation behavior needs to be adjusted according to proxemics rules, so it does not disturb people. The study of robot navigation in the presence of people is called Human-Aware Navigation (HAN).

Most approaches to HAN in the literature use only sensors on-board the robot. Even though those sensors bring the advantage of context-independence, it is useful to have other external sensors, which can add more information about the environment, not only in terms of coverage space but also in terms of precision on the estimation of the person's position. Hence, in this paper, in addition to the on-board camera, external cameras were mounted on the ceiling and used for the PD. This setup ensures a broader perception of the environment, capable of seeing both humans and robots at the same time.

All authors are with the Institute for Systems and Robotics (LARSyS), Instituto Superior Técnico, Universidade de Lisboa, Torre Norte - 7 Piso Av.Rovisco Pais, 1 1049-001 Lisboa, Portugal.

In this paper, we propose a novel solution for HAN resorting to multiple cameras (onboard and offboard), for people state estimation, using a novel deep learning technique. A cost-map is computed by combining several constraints associated with HAN. Each time the robot receives a new goal, it computes a path on that cost-map. When compared with state-of-the-art approaches, the main contributions presented in this paper are: a novel real-time technique for people detection; and a standardization of human-aware constraints.

The solution was tested in simulated and in realistic scenarios. The results show that the proposed solution fulfills the aforementioned goals. This paper proposes novel extensions of the authors' previous works, namely [1,2]. More specifically, in [1] HAN constraints are introduced, and in [2] a real-time PD algorithm is proposed.

A. State-of-the-Art

To conclude the introduction and to better understand the contributions presented in the paper, we review the related work on both PD and HAN, respectively.

1) PD: One of the main goals in our framework is to achieve an accurate and fast Pedestrian Detection algorithm. This has been one of the major topics addressed in the computer vision community, e.g. [3,4]. Classically, the PD problem has been addressed by using conventional handcrafted features (e.g., image gradient, HOG, wavelets, etc) that have plateaued in recent years. However, new achievements using deep compositional architectures, namely, Convolutional Neural Networks, have gained attention and have greatly advanced the performance of the state-of-the-art, concerning not only image classification but also in localization and detection [5]. As such, the application of deep learning to PD arises as a natural forthcoming step. Indeed, the advantage of deep learning based architectures [6] resides in the high-level features produced by the top layers of the model that allow to boost the classification results, compared to the performance produced by hand-crafted features [7].

A shortcoming is the training process for CNNs, which requires large amounts of annotated samples to avoid overfitting. This is, however, an apparent limitation since this issue has been tackled with transfer learning, which allows to retrain, via fine-tuning, publicly available models (e.g., models that have been trained with large annotated databases for other problems), using smaller datasets belonging to the target task [8].

To obtain real-time performance, faster PD methods based on handcrafted features (e.g., ACF [9]) can be combined with CNN models fine-tuned with PD datasets, to further improve the accuracy and eliminate false positives [10].

2) HAN: Regarding the HAN, in 2007 Sisbot et al. [11] proposed the HAN planner. Their work is focused on human comfort, which is addressed by three criteria: preventing personal space invasions; navigating in the humans' field of view (FOV); and preventing sudden appearances in the FOV of humans. Those criteria are modeled as cost functions in a 2D cost-map and path planning is performed with A^*. Even though the HAN planner accounts for replanning if people move, it does not adapt their personal space during the motion. With that in mind, two extensions to HAN were proposed:

- A prediction cost function which, by increasing the cost in front of a moving human, decreases the probability of the robot entering that area [12]; and
- The concept of compatible paths, which says that two paths are compatible if both agents can follow their paths (reaching the goal position), without any deadlocks [13].

An alternative approach was proposed by Kirby [14], which differs from HAN planner on the considered constraints and their formulation. Instead of focusing simply on human comfort, constraints concerning social rules (e.g., navigate on the right side of narrow passages) and low-level human navigation behavior (e.g., face direction of movement) are also taken into account. Another important issue related with human comfort, in a social context, is the interference with humans interacting with other humans and/or objects. This issue is tackled in [15] where, besides considering proxemics and the back space of a person, a constraint is included to model the space between interacting entities. Other important work was presented in [16]. The authors presented a framework for planning a smooth path through a set of milestones. Those are added, deleted, and/or modified, based on the static and dynamic components of the environment.

More recently, Kruse et al. [17] defined the three goals for Human-Aware Navigation: human comfort (e.g., space that people keep from each other in different contexts, known as the theory of proxemics [18], and velocity that robots navigate close to humans [19]); respect social rules; and mimic low-level human behavior.

B. Outline of the Paper

This paper is organized as follows. Section II introduces the main stages of the proposed framework. Section II-A describes the PD methodology. Section II-A1 details the CNN architecture used for PD. Section II-B describes the tracking procedure, i.e. how the above detections are associated between frames. Section III is related with how the deep learning methodology is integrated in the navigation setup. To accomplish this, the CNN training must be performed (Section III-A), as well as the adaptation of the CNN (Section III-B). Experimental evaluation is conducted in Section IV, in which, we evaluate the performance of the proposed PD methodology in the INRIA dataset (Section IV-A) and in two real scenarios comprising the "corridor" and "Mbot" sequences (Section IV-B). Section V presents the HAN constraints that result from the PD using one or more camera systems, and that will be used in the proposed framework. In Section VI we evaluate the HAN constraints and in Section VII we present the results with the complete framework (PD + HAN). Section VIII concludes the paper.

II. Vision-Based People Detection and Tracking through Deep Learning

For PD, this paper follows the strategy mentioned in the previous section, providing the following contributions. First, we adopt a very deep learning based approach, by using pre-trained models. Second, we are able to drastically decrease the computational burden that comes from a sliding window exhaustive search. To accomplish this, we cascade the ACF (non-deep) detector [9] with a CNN[1]. This cascade strategy is twofold: first, it provides a selective search approach that significantly improves the computational efficiency, since only the output proposals of the ACF are taken into account; then, by cascading with the CNN, we are able to boost the performance of the ACF detector (i.e., improve the classification accuracy of the ACF proposals by reducing the number of false positives). The proposed methodology uses the RGB feature map from the datasets. Fig. 1 illustrates the proposed approach for the PD task.

A. Methodology for PD

In this section we formalize the adopted methodology for PD. First, let us consider that we have available the following training set $\mathscr{D} = \{(\mathbf{x},\mathbf{y})_i\}_{i=1}^{|\mathscr{D}|}$, where $\mathbf{x}$ denotes the input image with $\mathbf{x}:\Omega \rightarrow \mathbb{R}^3$ with Ω denoting the image lattice[2] of size $w \times h \times d$, with $d = 3$; the class label is defined in $\mathbf{y} \in \mathscr{Y} = \{0,1\}^C$ that denotes the (absence) presence of the pedestrian in the ith image $\mathbf{x}_i$ (i.e., $C = 2$). The training dataset $\mathscr{D}$ is the input for the `ACF` detector. For each input image $\mathbf{x}_i$, a detector (e.g., `ACF` or `LDCF` [21]) is used to provide the candidate windows, or the proposals, along with the scores (confidences). This can be formalized as the following output set $\mathscr{O} = \{(\mathbf{x}(\mathscr{B}),\mathscr{S})_i\}_{i=1}^{|\mathscr{O}|}$. In this set, $\mathscr{B} = \{b_k\}_{k=1}^{|\mathscr{B}|}$ represents the set of bounding boxes coordinates, with $b_k = [x_k, y_k, w_k, h_k] \in \mathbb{R}^4$ denoting the top-left point and width and height enclosing (or not) the pedestrian; we denote as $\mathbf{x}(\mathscr{B})$, the content (i.e., the proposals) of the image delimited by the bounding boxes $\mathscr{B}$; $\mathscr{S} = \{s_k\}_{k=1}^{|\mathscr{S}|}$ are the `ACF` detector confidence scores assigned to the proposals $\mathbf{x}(\mathscr{B})$.

As mentioned in [8], the generalization ability of the CNN can be boosted resorting to pre-trained models, instead of using random initialization. Therefore, we use the recently proposed VGG CNN model [22][3], pre-trained with Imagenet [5]. Formally, we have a dataset to pre-train the CNN, i.e., $\widetilde{\mathscr{D}} = \{(\widetilde{\mathbf{x}},\widetilde{\mathbf{y}})_n\}_{n=1}^{|\widetilde{\mathscr{D}}|}$, with $\widetilde{\mathbf{x}}:\Omega \rightarrow \mathbb{R}^3$ and $\widetilde{\mathbf{y}} \in \widetilde{\mathscr{Y}} = \{0,1\}^{\widetilde{C}}$, where $\widetilde{C}$ is the number of classes in the pre-trained model (in the Imagenet case, it is 1000).

[1] Recall that there exist several non-deep detectors, such as Regionlets [20] or LDCF [21]. We choose the ACF detector since it is fast.

[2] In this paper, the RGB feature map is considered for the image $\mathbf{x}$.

[3] Details are also available at: http://www.robots.ox.ac.uk/~vgg/research/very_deep/.

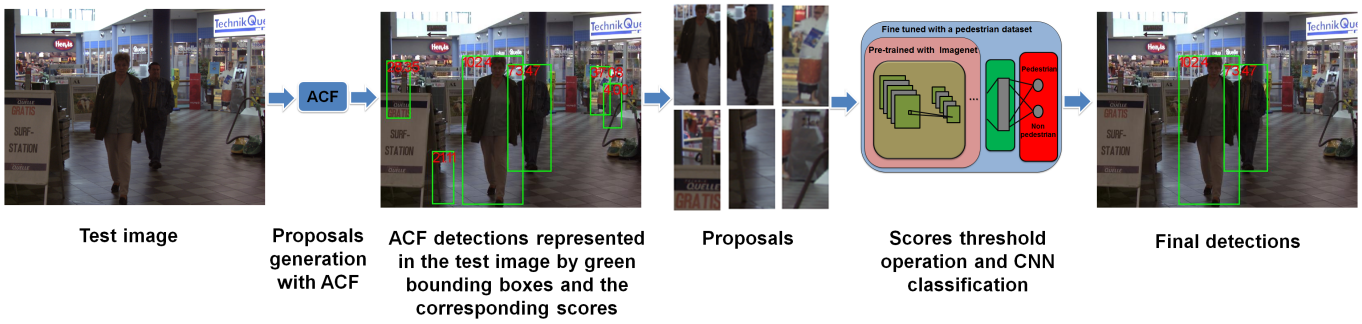

Fig. 1. Illustration of the proposed methodology, cascading the ACF (non-deep) detector and the deep CNN. First, the ACF detector performs selective search identifying promising image regions that contains pedestrians, generating pedestrian proposals (see the green rectangles in the third image of the figure). The candidate proposals, in the RGB feature map, are forwarded through the CNN to be more accurately classified (see text).

1) CNN model: We now formalize the main ingredients of the CNN architecture. Basically, this type of deep networks comprises several processing stages. Each stage is characterized by having two types of layers, namely: a convolutional layer containing a non-linear activation function, and a non-linear subsampling layer. In the former, a convolutional filter is applied to the input. In the latter, the size reduction of the input is achieved. These two stages are typically followed by several fully connected layers, and a multinomial logistic regression layer (see details in [7]). Formally, the CNN is defined as:

$$f(\mathbf{v},\theta^{(1)}) = \mathbf{v}^{\star} = f_{\text{out}} \circ f_L \circ ... \circ f_2 \circ f_1(\mathbf{v}^{(0)}), \tag{1}$$

where $\mathbf{v}^{(0)} = \mathbf{v}$ is the input data, $\circ$ denotes the composition operator, and $\theta^{(1)}$ represents the CNN parameters, i.e., the weights and biases. The output of the CNN mentioned in (1), can be seen as an approximation of the input data (represented by $\mathbf{v}^{\star}$). Each network layer contains a set of filters, where each filter is formally defined as:

$$\mathbf{v}^{(k)} = f_k(\mathbf{v}^{(k-1)}) = \sigma(\mathbf{W}_k(i,j)^{\top}\mathbf{v}^{(k-1)} + \beta_k), \tag{2}$$

where $\sigma(.)$ represents the non-linearity (e.g. the Rectified Linear Unit, for more details see [7]) and where the convolutional filters are represented by the weight matrix $\mathbf{W}_k$ and the bias vector β_k. The nonlinear subsampling layers are defined by $\mathbf{v}^{(k)} = \downarrow \mathbf{v}^{(k-1)}$, where $\downarrow$ denotes a subsampling function that pools (using the mean or max functions) the values from a region of the input data. The fully connected layers comprise a convolution as defined in (2), but the region to be processed is the entire input and the output locations are associated with individual filters. The multinomial logistic regression layer computes the probability of the i^{th} class using the features $\mathbf{v}^L$ from the L^{th} layer with the soft-max function $\mathbf{y}(i) = \frac{\exp(\mathbf{v}^L(i))}{\sum_j \exp(\mathbf{v}^L(j))}$. In our case, (1) is written as (similarly for (2)):

$$f(\mathbf{x}(\mathscr{B}),\theta^{(1)}) = \mathbf{x}(\mathscr{B})^{\star} = f_{\text{out}} \circ f_L \circ ... \circ f_2 \circ f_1(\mathbf{x}(\mathscr{B}^{(0)})), \tag{3}$$

where the inputs are the proposals (i.e., the image content, in the RGB feature map, delimited by the bounding boxes), here denoted as $\mathbf{x}(\mathscr{B}^{(0)})$ (see 4$^{\text{th}}$ image in Fig. 1). The main idea is to take the proposals $\mathbf{x}(\mathscr{B})$, that will be processed by the CNN, and produce a classification probability that a given proposal contains a pedestrian. The proposals classified as non pedestrians are discarded, allowing to eliminate false positives. The ones regarded as pedestrians are kept, including the original ACF detector score. In our case, the CNN prediction output can formally be represented by:

$$f(\mathbf{x}(\mathscr{B}),\theta) = y^{\star} \tag{4}$$

which is trained using the binary cross entropy loss over the training set indexed by i, as follows:

$$\mathscr{L} = \frac{1}{|\mathscr{D}|}\sum_{i=1}^{|\mathscr{D}|} -y(i)\times\log(y^{\star}(i)) - (1-y(i))\times\log(1-y^{\star}(i)) \tag{5}$$

Let the pre-trained CNN be represented by the model $\widetilde{\mathbf{y}} = f(\widetilde{\mathbf{x}},\widetilde{\theta})$, with $\widetilde{\theta} = [\widetilde{\theta}_{\text{cn}},\widetilde{\theta}_{\text{fc}},\widetilde{\theta}_{\text{lr}}]$. The process of pre-training a CNN is defined by the following three steps:

1) Training M_1 stages of convolutional and non-linear subsampling layers, that are represented by the parameters $\widetilde{\theta}_{\text{cn}}$; then
2) Training M_2 fully connected layers, represented by the parameters $\widetilde{\theta}_{\text{fc}}$; and
3) Training one multinomial logistic regression layer (with parameters $\widetilde{\theta}_{\text{lr}}$), by minimizing the cross-entropy loss function [7] over the dataset $\widetilde{\mathscr{D}}$.

It is worth to mention that, transferring a large number of pre-trained layers and fine-tuning the CNN is the key to achieve the best classification results in transfer learning problems [8]. Following this strategy, we first take the M_1 layers to initialize a new model (see [8]). Since we have changed the CNN input size to reduce the computational expense, the M_2 layers were randomly initialized from a Gaussian distribution, so that the dimensions are compatible and inference is possible. Finally, we introduce a new binomial logistic regression layer, with parameters θ_{lr} (randomly initialized from a Gaussian distribution) adapted for two classes (pedestrian and non-pedestrian). Afterwards, we fine-tune the CNN model by minimizing the cross-entropy loss function in (5), using the pedestrian training set $\mathscr{D}$.

B. Tracking

Taking the position measurements from the people detection scheme (described in the previous subsection), the goal of

the tracking phase is to associate detections between frames and to estimate the direction of the person's velocity. For that purpose, we use a simple Kalman filter. Moreover, since humans tend to walk at a constant velocity [23], a constant velocity model was assumed.

For each iteration, there are four possibilities (note that we are assuming that only one person is being tracked):

1) A person is not being tracked;
2) A person is being tracked and a new detection was received;
3) A person is being tracked and more than one detection was received;
4) A person is being tracked and no detection was received.

In the first case, if there is only one detection, the tracker is initialized. Otherwise (none or more than one detection was sent), that message is discarded. For the second case, the filter predicts a new state and then corrects it based on the measurement. If there is more than one detection, the measurement closest to the last estimated state is used to correct the filter's prediction. Finally, if there are no detections, the estimated state is the prediction from the filter, without correction.

III. Material and methods

This section addresses the implementation details of the deep learning methodology adopted, to be integrated in the navigation setup. First, in Sec. III-A we describe the experimental setup used to train the deep CNN. Then, Sec. III-B addresses the adaptation of the pre-trained model to the PD task. The outcome of these sections, i.e. the final CNN architecture fine-tuned for PD, is subsequently used for testing purposes, as described in Sec. IV.

A. Training CNN in the INRIA dataset

To train the CNN model we use the INRIA dataset, which is a common benchmark used for research work in the detection of pedestrians in images[4]. This dataset comprises 1832 training images, from which 1218 are negative images (i.e., not containing pedestrians) and 614 are positive images (i.e., containing pedestrians). There are 288 test images. In our experimental setup, we have to build not only the training set but also the validation set to fine-tune the pre-trained CNN model. To obtain the positive set, we use the ground truth positive training bounding boxes (i.e., proposals corresponding to the image content delimited by them) $\mathscr{B}_{\text{pos}} = 1237$. Data augmentation is then performed for the positive samples using the following two steps:

1) Horizontal flipping over the set $\mathscr{B}_{\text{pos}}$, resulting in a new set, $\mathscr{B}^{(1)}_{\text{pos}} = 2474$ (including also $\mathscr{B}_{\text{pos}}$); and
2) Random deformations (including translation and scale), in the range $R =]0,5[$ pixels applied on the previous set, $\mathscr{B}^{(1)}_{\text{pos}}$. This allows us to obtain a new set, $\mathscr{B}^{(2)}_{\text{pos}} = 4948$.

To build the negative set $\mathscr{B}_{\text{neg}}$, we extract negative windows (i.e., proposals) from the negative images using the strategy mentioned in [10] (i.e., employing a non-fully trained version of LDCF). As a result, a set of 12552 negative windows is obtained, by defining an upper-bound of 18 negative proposals per image. This strategy allows to acquire a total of 17500 proposals, from which 15751 are used for training (90%) and 1749 for validation (10%).

[4]More details can be found at: http://pascal.inrialpes.fr/data/human/.

B. Adaptation of the pre-trained CNN model

As mentioned in Sec. II-A, we use a pre-trained CNN model. As such, some modifications/adaptations are necessary. In the following, we describe the required steps, by first detailing the model used and, then, describing the modifications needed.

1) Pre-trained model used: For the pre-trained network, we select the VGG Very Deep 16 architecture (VGG-VD16), more specifically, the "D" configuration in [22][5].

This CNN architecture receives $224 \times 224 \times 3$ input images. The network comprises 13 convolutional layers, three fully connected layers and a multinomial logistic regression layer (see Sec. II-A1). There are five max-pooling operations (i.e., non-linear subsampling), which operate on a 2x2 region, reducing the input by a factor of two. These operations appear after layers: 2, 4, 7, 10 and 13. The Rectified Linear Unit (ReLU) is selected to be the non-linearity, applied to the convolutional and fully connected layers. Furthermore, the size of the receptive field is the same for all convolutional layers (i.e., 3x3). However, the number of filters is different as described below:

- There are 64 filters in each of the layers 1 and 2;
- There are 128 filters in each of the layers 3 and 4;
- There are 256 filters in each of the layers 5, 6 and 7;
- There are 512 filters in each of the layers from 8 to 13;
- There are 4096 filters in each of the layers 14 and 15;
- There are 1000 filters in layer 16. Each of these filters is associated with a different ILSVRC [5] class. After this layer, there is a soft-max function.

The pre-training of the mentioned model was performed with Imagenet [5], having 1K visual classes, 1.2M training, 50K validation and 100K test images.

2) Adaptation of the deep network: In order to decrease the computational effort required by the VGG-VD16 model, the expected input size for the CNN was reduced from $224 \times 224 \times 3$ (i.e., the original) to $64 \times 64 \times 3$. Nevertheless, due to this change in the input size, inference is not possible after the first fully connected layer. Moreover, the PD task requires the use of only two classes in order to represent the presence or absence of a pedestrian. Therefore, the layer 16 and the soft-max must be adjusted accordingly. In order to solve the aforementioned issues, we choose to randomly initialize the parameters of the three fully connected layers. More specifically, they are obtained from a Gaussian distribution, with zero mean and variance $\sigma^2 = 0.01$. The changed network is fine-tuned using the previously mentioned (in Sec. III-A) positive and negative sets, obtained from the INRIA pedestrian dataset. The process of pre-training the network acts

[5]See additional details at: http://www.robots.ox.ac.uk/~vgg/research/very_deep/.

as a regularization procedure, similar to data augmentation. Regarding the fine-tuning hyperparameters, we use $\varepsilon = 10$ epochs, a minibatch of $\xi = 100$ samples, a learning rate of $\beta = 0.001$, and a momentum of $\mu = 0.9$. No special effort was dedicated to fine-tune the mentioned hyperparameters.

During test, we run the ACF detector in the 288 INRIA test images in order to obtain the proposals (i.e., regions potentially containing pedestrians). Then, we run the proposals through the CNN in order to classify them as pedestrians or non-pedestrians.

All the experiments will be detailed next (see Sec. IV and Tables I, II and III) and were obtained with Matlab, running on CPU mode on 2.50 GHz Intel Core i7-4710 HQ with 12 GB of RAM and 64 bit architecture. The Piotr's Computer Vision Matlab Toolbox [24] (2014, version 3.40) was employed to execute the ACF method, and to perform the performance evaluation. Regarding the implementation of the CNN framework, the MatConvNet toolbox [25] was utilized.

IV. Evaluation of the People Detection

This section provides the testing results for the evaluation of the proposed PD method. The section is divided into two parts:

- In Sec. IV-A, we describe the performance evaluation, concerning both accuracy and runtime figures of the PD method, when the INRIA dataset is used; and
- In Sec. IV-B, we evaluate the performance of the PD method in real scenarios comprising two sequences (termed herein as "corridor" and "Mbot"). Special attention is given to the details of how we achieve real-time requirements in the detection.

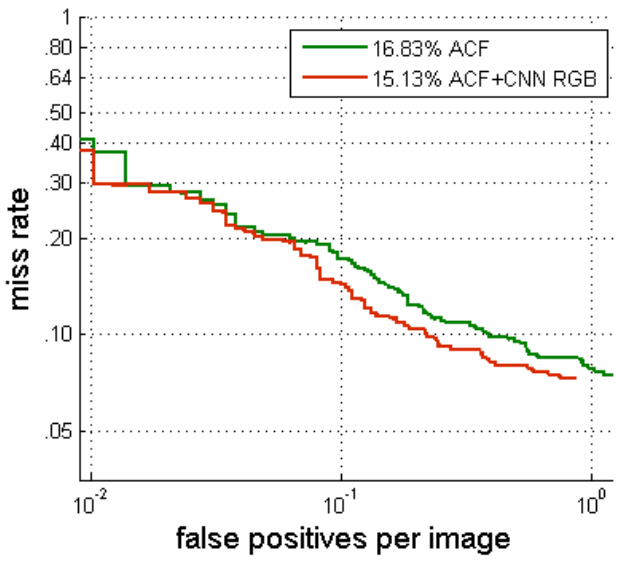

Fig. 2. Graph of the miss rate versus false positives per image for the INRIA dataset, using the ACF detector alone and the proposed approach ACF+CNN RGB. The values in the box represent the log average miss rate for each method.

TABLE I
NUMBER OF TRUE POSITIVES (TP), FALSE POSITIVES (FP) AND FALSE NEGATIVES (FN) FOR THE ACF PROPOSALS ON THE INRIA DATASET, AND THE POSTERIOR CNN CLASSIFICATION. THE RGB FEATURE MAP WAS USED FOR THE FINE-TUNING.

Dataset	Metrics	ACF Proposals	ACF+CNN RGB
INRIA	TP	551	546
	FP	1284	249
	FN	38	43

TABLE II
RUNNING TIME FIGURES AND MISS RATE (MR) USING THE PRE-TRAINED MODEL, OBTAINED FOR INRIA DATASET AND CORRESPONDING LOG AVERAGE MISS RATE

Dataset	Pre-trained VGG-VD16
INRIA	MR ACF+VGG proposals = 15.13% ACF train time = 5.1 hrs. CNN test time = 67.3 sec.

A. Performance evaluation on the INRIA dataset

To test the accuracy of the PD module, to be integrated in the navigation setup, we first assess its performance in the INRIA dataset. The adopted evaluation metric is the log average miss rate, as proposed in [3]. This metric is obtained from nine values in the False Positives Per Image (FPPI) interval $[10^{-2}, 10^{0}]$.

To demonstrate the effectiveness of the proposed cascade of a non-deep detector (ACF) with the CNN architecture, we also present the results of the ACF alone and the results of the cascade to notice the improvement achieved. From the experiments conducted, we achieve a log average miss rate of 16.83% for the ACF detector alone. Cascading the ACF with the CNN we are able to reach 15.13%.[6] To enrich the evaluation of the algorithm's performance (and the ACF baseline comparison) at several FPPI operating conditions, Fig. 2 displays the graph of the miss rate versus FPPI (more information is available in [3]).

Table I shows the number of true positives (TP), false positives (FP) and false negatives (FN), before and after the use of the CNN. The number of FP is significantly reduced, while maintaining most of the TP. This indicates that the CNN is successfully discarding FP, allowing to reach performance improvements (and justifying the observed gain).

An important concern when building the PD algorithm, is to verify and ensure that the runtime figures achieve the real-time requirements. Table II shows the CNN training and testing running times obtained for the INRIA dataset. From the results, we conclude that the system is able to perform the PD up to 4.2 frames per second (FPS) [7], which should be improved. This concern will be addressed in the next section.

[6]The approach proposed herein, as shown in Fig. 1, is general and can be applied to any other non-deep detector (e.g. Regionlets [20], LDCF [21] or Spatial Pooling [26] detectors).

[7]This metric refers to the CNN alone, not considering the ACF runtime for the proposals generation.

B. Performance evaluation on real scenarios

To perform experiments in real scenarios, we acquired two indoor datasets to evaluate the PD task. The main objective is to reach the real-time detection requirement, i.e. to improve the runtime figures shown in Tab. II, while achieving as much accuracy as possible. Two datasets are considered in these experiments:

- The "corridor" dataset, containing 5556 images; and
- The "Mbot" dataset, containing 3966 images.

The size of the frame (i.e., the image) is 480×640 for both of the sequences. Some of the results obtained in the two datasets are shown in Fig. 3, where each bounding box has a score, mentioning the confidence of containing a pedestrian.

The total time to perform PD in the dataset "corridor" (5556 frames) and "Mbot" (3966 frames), is roughly 707.27 minutes and 839 minutes, respectively. The running time figures per frame are shown in Tab. III (top). In this table (top), the "baseline" refers to the final PD detector as detailed in Sec. III-B. The algorithm reaches 7.85 FPS and 4.84 FPS, for the "corridor" and the "Mbot" sequences, respectively. Although these results are quite reasonable, it is possible to improve them. Basically, two strategies can be used to improve the above results: (a) to reduce the original images size (480×640), and (b) to discard ACF proposals with confidence score below a certain threshold. In this work, we opt to perform the second one. Notice that, reducing the size of the images may jeopardize the quality of the detections. Furthermore, the confidence scores outputted by the ACF detector constitute a relevant indicator to filter the proposals. This can be achieved by applying a threshold over the proposals, where only the ones above a certain score are kept and are further processed by the CNN. From our experiments a threshold value of 40 was set for the confidence score. This threshold is in accordance with the value suggested in [27], in which it is proposed a thresholding technique based on upper and lower bounds on the confidence scores of the ACF. Another important remark is that, the gain in speed results from the fact that the CNN only has to classify a smaller portion of the ACF proposals, instead of all of them. Therefore, the easier false positives should be discarded by the threshold operation, while the harder false positives should be discarded by the CNN. The threshold value controls the trade-off between potential accuracy loss and speed gain. By selecting the later scenario, we are able to improve the runtime of the overall detector (i.e., ACF+CNN), when compared with the results obtained for the baseline. These results are shown in Table III (bottom), where it can be seen that a frame rate of approximately 10 FPS is now obtained, which is suited for real-time applications.

TABLE III
RUNNING TIME FIGURES FOR THE BASELINE OF THE ARCHITECTURE (TOP) AND PERFORMING A THRESHOLD OPERATION OVER THE ACF PROPOSALS (BOTTOM).

Dataset	Data seq. 1 (corridor)	Data seq. 2 (Mbot)
Baseline	Total time = 0.1273 sec. ACF time = 0.0326 sec. CNN time = 0.0947 sec. **frame rate = 7.85** FPS	Total time = 0.2066 sec. ACF time = 0.0367 sec. CNN time = 0.17 sec. **frame rate = 4.84** FPS
Threshold	Total time = 0.0961 sec. ACF time = 0.0333 sec. CNN time = 0.0628 sec. **frame rate = 10.41** FPS	Total time = 0.1026 sec. ACF time = 0.0381 sec. CNN time = 0.0645 sec. **frame rate = 9.74** FPS

V. HUMAN-AWARE NAVIGATION

From the method described in the previous section, we can define bounding boxes on the images that represent pedestrians. One can use several cameras and this PD in each camera. From the assumption that we have a set of images (a single image can also be used), the goal of this section is to first fuse the information given from the different camera sensors (if there is more than one) and then define the HAN constraints that can be later included in any conventional path planer for the complete HAN.

After computing the position of the pedestrian (or pedestrians) on the image planes (method described in the previous sections) and converting its position into the world coordinate systems, including the estimation of the pedestrian's velocity (the methods used in this paper for this module are shown in Sec. V-A), the main goal of this section is to define the robot's path in a scenario with humans that may interact with it.

To be as Human-Aware as possible, three goals were considered, namely: human comfort; social rules respect; and naturalness. To fulfill these requirements, the following constraints were taken into account:

1) Take least effort path (naturalness) – Sec. V-B;
2) Keep a distance from static obstacles (naturalness) – Sec. V-B;
3) Respect personal spaces (human comfort) – Sec. V-C1;
4) Avoid navigating behind sitting humans (human comfort) – Sec. V-C2;
5) Do not interfere with human-object interactions (human comfort) – V-C3; and
6) Overtake people by the left (social rule) – Sec. V-C4;

The first two constraints are related to navigations problems (the method used in this paper is described in Sec. V-B), whereas the remaining four constraints are about the HAN (more details about each of these constraints are shown in Sec. V-C).

A. People Tracking in the World

As explained in the previous sections, the PD module returns bounding boxes representing people in the scene. Those are sent to this module (People Tracking module), which uses an array of Kalman Filters to track people (one for each person). It performs the following steps:

1) Project the middle point, between the bottom left and right corners of each bounding box, to the ground plane, in world coordinates;
2) Predict the new state of each track;
3) Associate the measurements to existing tracks and determine if new tracks need to be created or removed;
4) Update the prediction with the respective measurement; and
5) Create and/or remove signaled tracks.

(a) Pedestrian detection in the "corridor" sequence.

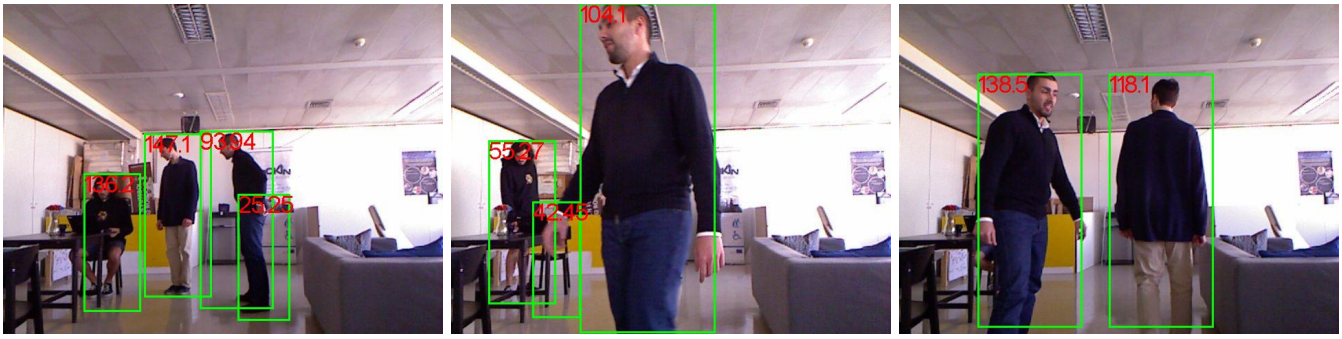

(b) Pedestrian detection in the "Mbot" sequence.

Fig. 3. In real scenarios, there are cameras mounted on the ceiling and on the robot. To test our PD, we used two sequences of images acquired from both possible real scenario camera locations. In Figs. (a) and (b) are shown three images of the "corridor" and "Mbot'" sequences, respectively.

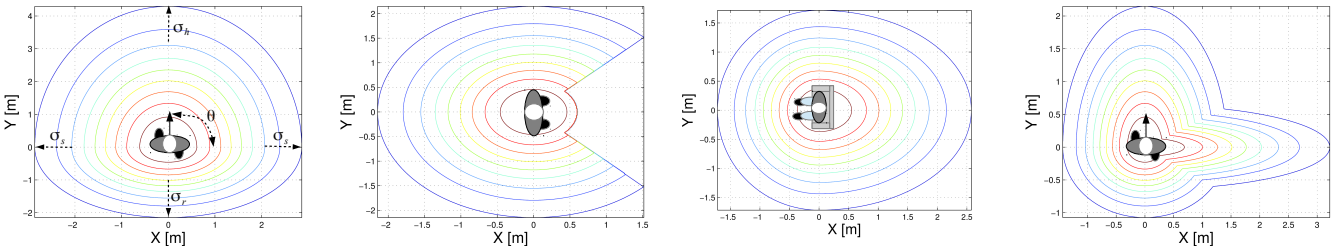

(a) Personal space (b) Object hand over (c) Seated person (d) Walking person

Fig. 4. Representation of cost functions associated with different people postures: Fig. (a) represents the cost function for the personal space of a person walking in the y direction at 1 [m/s]; Fig. (b) shows the cost function of a person standing, oriented in the x direction, during an object hand over; Fig. (c) represents the cost function for the case where a person is seated; and Fig. (d) shows the total cost function of a walking person, including the social rule of overtaking her by the left and her personal space.

The people tracking is performed in the world coordinate system, instead of the image plane. For this purpose, two assumptions are considered: i) people are standing or walking upright; ii) given i) a person feet will always be on the ground plane and, thus, the point, which represents a person feet, is on the line between the bottom left and right corners of the respective bounding box. Since the person is, most of the times, in the center of the bounding box, a good estimate for the position of the feet is in the middle point of the considered line segment. The projection is performed by transforming the selected point with an homography computed a priori (transformation from the image plane into the floor plane). These positions will be used as measurements in Kalman Filters (one for each person), which will perform the tracking. A constant acceleration motion model was considered for the prediction step of the filter.

The association is performed with a Greedy method. The greedy method uses all the distances between tracks and measurements. It will associate a measurement to the closer track automatically. When the association is complete, there is a condition that checks if the assignment distance is higher than a threshold. When that is the case, the track and measurement are unassigned. After the assignment, there are three cases to be considered: the assignment exists; a track is not assigned; and/or a measurement is not assigned. If there is a successful assignment, it proceeds to the correction step of the filter, with the assigned measurement. When a track is not assigned, there are two cases: the track increases the inactivity flag or, if the inactivity threshold was reached, the track is deleted. If a measurement is not assigned, after some frames, a new Kalman Filter is created for this measurement.

B. Path planner and obstacle avoidance

The first HAN constraint (see the complete list in Sec. V) is addressed by the path planner. In this work the A^* algorithm [28] was used, ensuring a minimum cost path as long as the heuristic is admissible. The total cost of a node is given by the sum of the cost of reaching that node, with the heuristic cost. The latter was considered to be the Euclidean distance from the node to the goal position. Since the environment is dynamic (people may appear walking in the scene), the planner computes a path periodically.

The goal of the second constraint is to prevent the robot from passing too close to obstacles. This is solved by attributing a high cost to the area surrounding the obstacles, [29]. In the next subsection, we define a set of constraints, that will be included in the path planner, for the navigation to be Human-Aware.

C. Human-Aware Navigations constraints

The Human-Aware Constraints (defined in Sec. V) used in this work are based on previous state-of-the-art approaches. However, some of them are reformulated (namely constraints 4 and 5) in order to be better integrated in our approach and to standardize their formulation.

1) Personal Space Constraints: The third constraint in the list accounts for personal space. We consider three different situations: when a person is standing; walking; or seated. For the case of a person walking, we used the formulation proposed in [14], which the authors call asymmetric Gaussian:

$$g = \text{asymGauss}(e_x, e_y, \alpha, \beta_h, \beta_s, \beta_r), \tag{6}$$

where

- α – orientation of the function;
- β_h – variance in the α direction;
- β_r – variance in the $\alpha - \pi$ direction;
- β_s – variance in the $\alpha \pm \frac{\pi}{2}$ direction; and
- e_x and e_y – are the variables that define the space around the person.

A graphical representation of these parameters is shown in Fig. 4(a). Then, the personal space of a walking person was modelled as:

$$g_1 = \text{asymGauss}(e_x, e_y, \alpha_p, \widetilde{\beta}, \frac{2}{3}\widetilde{\beta}, \frac{1}{2}\widetilde{\beta}), \text{ where } \widetilde{\beta} = \max(v, 0.8), \tag{7}$$

where α_p is the person's orientation and v her speed. A graphical representation of the person walking along e_y–axis direction with a velocity of 1 [m/s] is presented in Fig. 4(a).

Regarding a walking person, it makes sense for the personal space in front to be larger than in the back (to ensure the robot does not pass in front of the person, decreasing the risk of collision). On the other hand, if a person is standing and we consider the personal space defined using the previous formulation, the robot may pass behind too close to the person, causing discomfort. Thus, for this case we suggest that the personal space to be modelled as a circular Gaussian:

$$g_2 = \exp\left(-\frac{(e_x - p_x)^2}{2\beta_{e_x}^2} - \frac{(e_y - p_y)^2}{2\beta_{e_y}^2}\right), \tag{8}$$

where (p_x, p_y) is the person position, β_{e_x} and β_{e_y} are the standard deviation in the e_x and e_y direction respectively. This formulation was also considered for a seated person.

If an object hand over is required, the robot should be able to enter the personal space, to be at "arm's length". However, the robot cannot be allowed to enter from a random direction, instead it should only be allowed to approach a person from the front, [30]. Thus, our solution is to open the region in front of the person 45 degrees, to a distance of 0.6[m]. Personal space, in a hand over scenario, is depicted in Figure 4(b).

2) Visibility Constraint: Constraint 4 concerns preventing discomfort from passing behind a seated person. We reformulated this problem with an asymmetric Gaussian:

$$g_3 = \text{asymGauss}(e_x, e_y, \alpha_p - \pi, 1.2, 0.8, 0.006). \tag{9}$$

3) Interaction Constraint: The fifth constraint prevents the robot from interfering with a person interacting with an object. It is represented by an interaction set modelled as a circle:

$$g_4 = \begin{cases} \alpha & \text{if} \quad (e_x - \widetilde{p}_x)^2 + (e_y - \widetilde{p}_y)^2 \leq r \\ 0 & \text{otherwise} \end{cases}, \tag{10}$$

where the middle position between the interacting entities are denoted as $(\widetilde{p}_x, \widetilde{p}_y)$, and the radius r is half the distance between the entities (only one-on-one interactions are considered). α is an importance factor, which varies from 0 to 1.

4) Overtake Constraint: Constraint 6 represents the social rule of overtaking people by the left (considered only for walking persons). This constraint is also represented using an asymmetric Gaussian:

$$g_5 = \text{asymGauss}(p_x, p_y, \alpha_p - \frac{\pi}{2}, 1.5, 0.3, 0.0075), \tag{11}$$

For the three possible postures of a person (standing, seated and walking), there are two where more than one cost function is applied and they must be combined. Since the main goal of the framework is to maximize the comfort of the humans, the cost functions are combined by taking the maximum cost value attributed to each point in space. The first case of multiple cost functions affecting the same space, is a seated person, whose personal space is given by:

$$g_6 = \max(g_2, g_3), \tag{12}$$

this cost function is depicted in Fig. 4(c). The second case concerns a walking person, where the personal space must be combined with the respective social rule (a person should be overtaken by the left):

$$g_7 = \max(g_1, g_5). \tag{13}$$

A graphical representation of this cost function is shown in Fig. 4(d).

VI. Evaluation of the Proposed Human-Aware Navigation Constraint

To evaluate the proposed constraints for HAN, four experiments were defined and tested in simulated environments. The proposed system was implemented as an extension of the ROS navigation stack, [29], and the cost functions, described in the previous section, were implemented as plug-ins to the

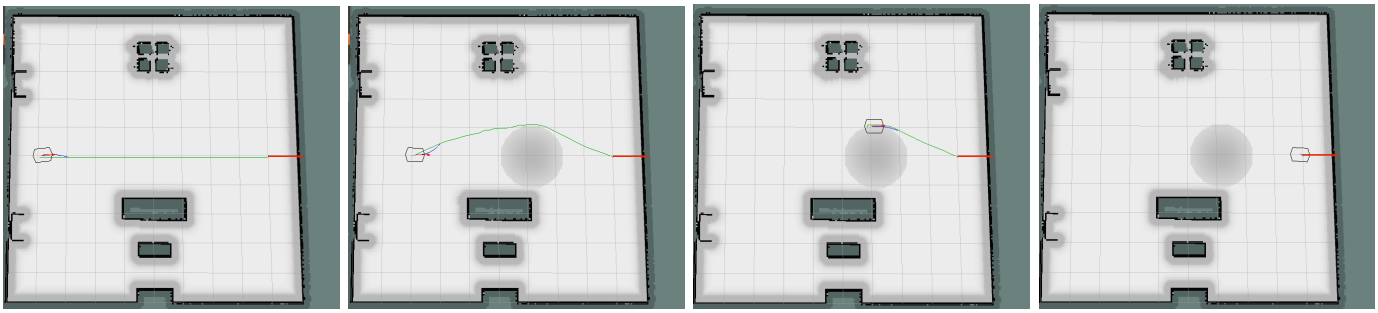

(a) The robot is navigating in free space and a standing person appears on the robot's path. The goal is to see if the robot is able to replan its path, without violating personal space.

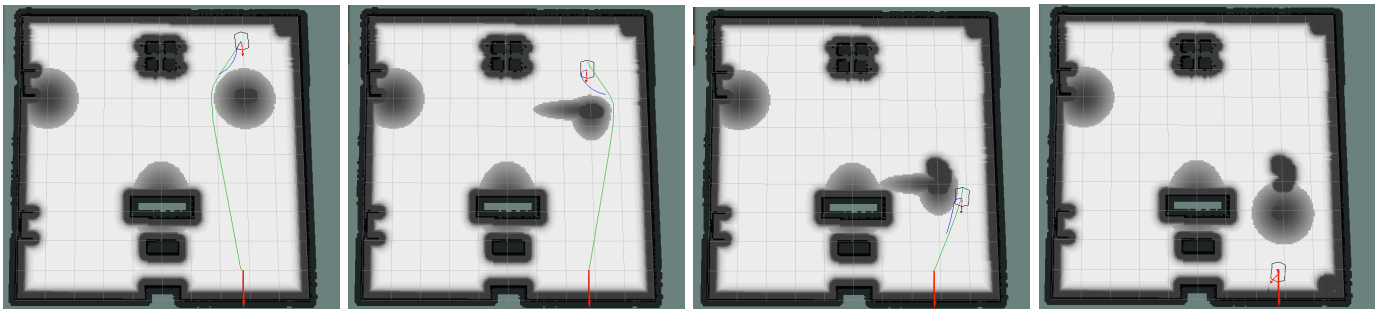

(b) The robot was placed behind a person and a goal position was defined. Then the person starts walking, the robot should replan a path through the person's left.

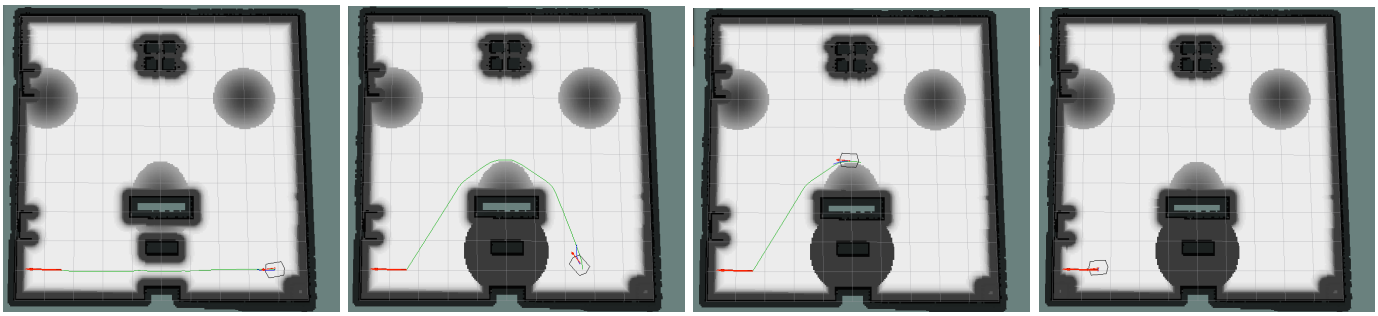

(c) A person is seated in front of a TV and a goal position is given to the robot. That goal will imply that the robot will pass between the person and TV. After that, TV is turned on and the robot should replan the path avoiding that area.

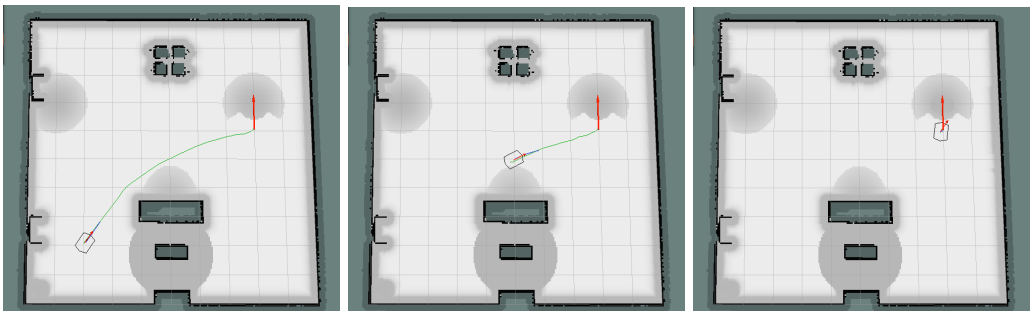

(d) The robot navigates towards a person, to hand over some object. The goal is to verify the modification of the personal space.

Fig. 5. Evaluation of the proposed navigation system using simulated environments. The environment was created using Gazebo and the results are shown in Rviz (ROS package). Figs. (a) (b) (c) and (d) show sequences of images representing experiments 1, 2, 3, and 4 respectively (more details regarding each of the experiments are given in the text). In all cases, it can be seen: the cost-map, current and goal position, path, and the trajectory.

cost-map layered structure, [31]. For the implementation of computer vision algorithms, we use the OpenCV framework [32].

The simulation environment was Gazebo. For the vision-based person tracking, a computer with an Intel Core i5-2430M, with 6GB of RAM (external CPU) was used. For navigation components (reconfiguring cost-maps, path planning, and trajectory execution), we used an Intel Core i7-3770T with 8GB of RAM (on-board CPU).

The experiments performed were:

Experiment 1: The robot is navigating in free space and a standing person appears on the robot's path. The goal is to see if the robot is able to replan its path, without violating personal space.

Experiment 2: The robot is navigating when encounters a slow walking person, which it must overtake. The goal is to verify if it respects constraints 3 and 6, Sec. V.

Experiment 3: A person is seated on a couch, watching TV, and the robot wants to go across the room. The goal is to test if the robot respects constraints 4 and 5, Sec. V.

Experiment 4: The robot navigates towards a person, to hand over some object. The goal is to verify the modification of the personal space, constraint 3, Sec. V.

The results for these experiments are shown in Fig. 5.

Next, we present experimental results using both simulated and realistic environments.

VII. Results of the Complete Framework

In this section, we evaluate the proposed framework (using both the proposed PD and the proposed HAN). For that purpose, we use a MBOT mobile platform [33], in a typical domestic indoor scenario, as shown in Fig. 6.

For the PD, for now, we are using a single camera to compute the person position (more specifically, the camera that gives the image of Fig. 6(c)). In the future, we are planning to use all the cameras shown in the setup, Fig. 6.

For the validation, we consider the following experiment (as a future work, more experiments are planned): a robot is navigating towards a goal position, and two people are standing in the environment. At some point, during the motion, a person starts walking, blocking the robot's path. At that moment, the robot must overtake the pedestrian according to the respective social rule (see Sec. V-C). The results are shown in Fig. 7.

Regarding the HAN, throughout the experiments the robot displayed a similar behaviour to the simulation in terms of trajectory execution in most cases. However, the parameters of the cost functions (7), (8), (9), and (11) needed to be adjusted. The values presented previously were derived empirically, taking into account: the values in the literature; the space restrictions of the real scenario, and our intuition of comfort distances.

VIII. Conclusions

This work addresses the problem of Human-Aware Navigation for a robot in a social context. For that purpose, in this paper we derive a robust and real-time solution for the PD, using deep learning. In addition, regarding HAN, we reformulate some of the respective constraints in order to have a standardization of human-aware constraints.

To validate our contributions, we first use the INRIA dataset to evaluate both the accuracy and runtime figures of the PD method. Then, to evaluate the performance using real data, we use two sequences of images (one was acquired using the robot's onboard camera and the second sequence was acquired using an external camera). The results show that the method is both robust and fast enough (up to 10 frames per second) for robot navigation applications. To evaluate the standardization of HAN constraints, we use a simulated environment.

In a realistic scenario, we test both modules in an example where the robot firstly finds two people standing and then, when it is moving towards the goal, a person appears, blocking the path. According to the social rule, the robot must overtake the person through the left. The results show that the robot has the correct behavior, which means that both PD and HAN are working properly. More experiments in realistic scenarios are expected, by fusing data from external and on-board cameras and experiments with different pedestrian behaviors.

References

[1] A. Mateus, P. Miraldo, P. U. Lima, and J. Sequeira, "Human-Aware Navigation using External Omnidirectional Cameras," *Iberian Robotics Conference*, 2015.

[2] D. Ribeiro, A. Mateus, J. C. Nascimento, and P. Miraldo, "A Real-Time Pedestrian Detector for Robot Navigation," *arXiv: Robotics*, 2016.

[3] P. Dollar, C. Wojec, B. Schiele, and P. Perona, "Pedestrian detection: An evaluation of the state of the art," *IEEE Trans. Pattern Analysis and Machine Intelligence*, 2012.

[4] R. Benenson, M. Omran, J. Hosang, and B. Schiele, *Computer Vision - ECCV 2014 Workshops: Zurich, Switzerland, September 6-7 and 12, 2014, Proceedings, Part II*. Cham: Springer International Publishing, 2014, ch. Ten Years of Pedestrian Detection, What Have We Learned?, pp. 613–627.

[5] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *Int'l J. Computer Vision*, 2015.

[6] I. Goodfellow, Y. Bengio, and A. Courville, "Deep learning," 2016, book in preparation for MIT Press.

[7] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Neural Information Processing Systems (NIPS)*, 2012.

[8] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" in *Neural Information Processing Systems (NIPS)*, 2014.

[9] P. Dollar, R. Appel, S. Belongie, and P. Perona, "Fast feature pyramids for object detection," *IEEE Trans. Pattern Analysis and Machine Intelligence*, 2014.

[10] D. Ribeiro, J. C. Nascimento, A. Bernardino, and G. Carneiro, "Improving the performance of pedestrian detectors using convolutional learning," *Pattern Recognition*, 2016.

[11] E. A. Sisbot, F. M.-U. Luis, R. Alami, and T. Simeon, "A Human Aware Mobile Robot Motion Planner," *IEEE Trans. Robotics*, 2007.

[12] T. Kruse, A. Kirsch, E. A. Sisbot, and R. Alami, "Exploiting Human Cooperation in Human-Centered Robot Navigation," *IEEE Int'l Symposium on Robot and Human Interactive Communication (RO-MAN)*, 2010.

[13] T. Kruse, P. Basili, S. Glasauer, and A. Kirsch, "Legible Robot Navigation in the Proximity of Moving Humans," *IEEE Workshop on Advanced Robotics and its Social Impacts (ARSO)*, 2012.

[14] R. Kirby, "Social Robot Navigation," Ph.D. dissertation, Robotics Institute, Carnegie Mellon University, 2010.

[15] L. Scandolo and T. Fraichard, "An Anthropomorphic Navigation Scheme for Dynamic Scenarios," *IEEE Proc. Int'l Conf. on Robotics and Automation (ICRA)*, 2011.

(a) Image of the robot platform used for the experimental results in realistic cenarios.

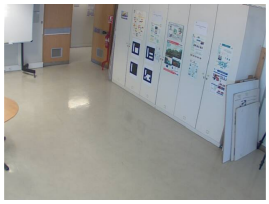

(b) Image of `Camera 1`, mounted on the ceiling.

(c) Image of `Camera 2`, mounted on the ceiling.

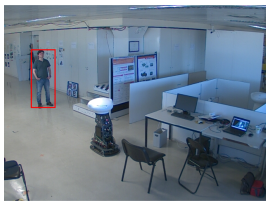

(d) Image of `Camera 3`, mounted on the ceiling.

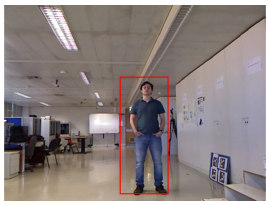

(e) On-board robot camera's image.

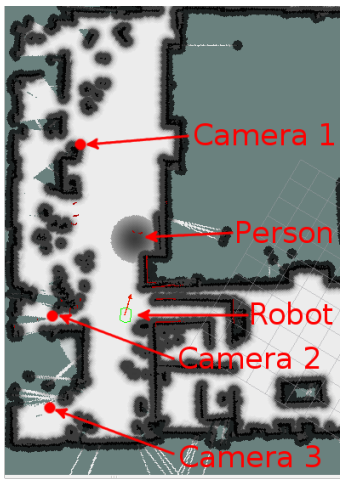

(f) Depiction environment in Rviz, showing: the robot's position, a pedestrian, and the positions of the cameras.

Fig. 6. Representation of setup used in the experiments in a realistic scenario. Fig. (a) shows the robot platform and Figs. (b), (c), (d), and (e) show the images of the cameras that will be used to detect the pedestrians (as it can be seen, in these images we already show the bounding boxes identifying a person in the environment). To conclude, Fig. (f) shows the environment (ROS Rviz package), with the position of all the cameras, the position of the robot and the pedestrian, with the respective HAN constraint (in this case the pedestrian was standing).

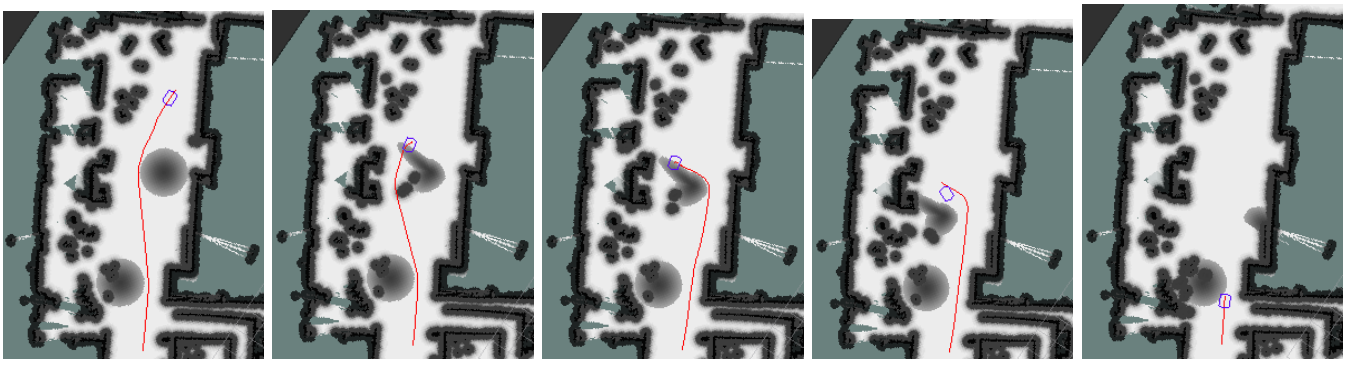

Fig. 7. Results of the real experiments in a realistic scenarios. A MBOT mobile robot [33] is used on a typical domestic indoor scenario. In this experiment, robot is navigating while a person (which firstly was stand) starts walking blocking the robot's path. At that moment, the robot must overtake the pedestrian according to the respective social rule that says that a robot must overtake a person through the left. The robot is shown as a blue rectangle while the path of the robot is shown as a red line.

[16] A. Pandey and R. Alami, "A Framework towards a Socially Aware Mobile Robot Motion in Human-Centered Dynamic Environment," *IEEE/RSJ Proc. Int'l Conf. Intelligent Robots and Systems (IROS)*, 2010.

[17] T. Kruse, A. K. Pandey, R. Alami, and A. Kirsch, "Human-aware robot navigation: A survey," *Robotics and Autonomous Systems*, 2013.

[18] E. Hall, *The Hidden Dimension: Man's Use of Space in Public and Private*. Bodley Head, 1969.

[19] D. Shi, E. G. Collins Jr., B. Goldiez, A. Donate, X. Liu, and D. Dunlap, "Human-Aware Robot Motion Planning with Velocity Constraints," *Int'l Symposium Collaborative Technologies and Systems (CTS)*, 2008.

[20] X. Wang, M. Yang, S. Zhu, and Y. Lin, "Regionlets for generic object detection," in *ICCV*, 2013.

[21] W. Nam, P.Dollar, and J. H. Han, "Local decorrelation for improved pedestrian detection," in *NIPS*, 2014.

[22] A. Z. K. Simonyan, "Very deep convolutional networks for large-scale image recognition," *ICLR*, 2015.

[23] S. Bitgood and S. Dukes, "Not Another Step! Economy of Movement and Pedestrian Choice Point Behavior in Shopping Malls," *Environment and Behavior*, 2006.

[24] P. Dollár, "Piotr's Computer Vision Matlab Toolbox (PMT)," http://vision.ucsd.edu/~pdollar/toolbox/doc/index.html.

[25] A. Vedaldi and K. Lenc, "MatConvNet – Convolutional Neural Networks for MATLAB," in *Proceeding of the ACM Int. Conf. on Multimedia*, 2015.

[26] S. Paisitkriangkrai, C. Shen, and A. van den Hengel, "Strengthening the effectiveness of pedestrian detection with spatially pooled features," in *ECCV*, 2014.

[27] A. Verma, R. Hebbalaguppe, L. Vig, S. Kumar, and E. Hassan, "Pedestrian detection via mixture of cnn experts and thresholded aggregated channel features," in *ICCV Workshop*, 2015.

[28] P. E. Hart, N. J. Nilsson, and B. Raphael, "A Formal Basis for the Heuristic Determination of Minimum Cost Paths," *IEEE Trans. Systems Science and Cybernetics*, vol. 4, no. 2, July 1968.

[29] E. Marder Eppstein, E. Berger, T. Foote, B. Gerkey, and K. Konolige, "The Office Marathon: Robust Navigation in an Indoor Office Environment," *IEEE Proc. Int'l Conf. Robotics and Automation (ICRA)*, 2010.

[30] K. L. Koay, E. A. Sisbot, D. S. Syrdal, M. L. Walters, K. Dautenhahn, and R. Alami, "Exploratory Study of a Robot Approaching a Person in the Context of Handing Over an Object." *AAAI spring symposium: multidisciplinary collaboration for socially assistive robotics*, 2007.

[31] D. V. Lu, D. Hershberger, and W. D. Smart, "Layered Costmaps for Context-Sensitive Navigation," *IEEE/RSJ Proc. Int'l Conf. Intelligent Robots and Systems (IROS)*, 2014.

[32] "OpenCV (Open Source Computer Vision Library)," http://opencv.org/about.html, accessed: 2016-04-26.

[33] J. Messias, R. Ventura, P. Lima, J. Sequeira, P. Alvito, C. Marques, and P. Carriço, "A Robotic Platform for Edutainment Activities in a Pediatric Hospital," in *IEEE Int'l Conf. Autonomous Robot Systems and*

Competitions (ICARSC), 2014.